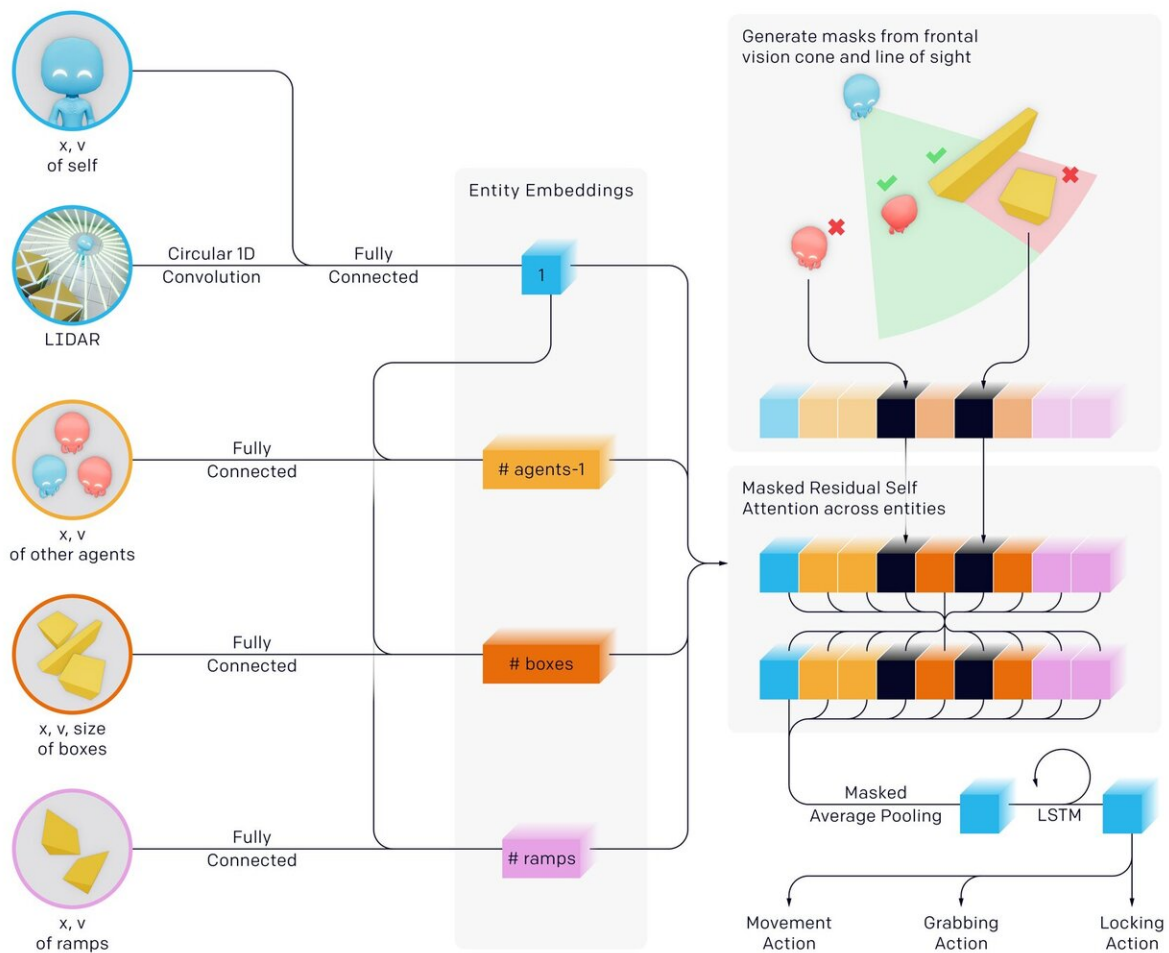


AI: Agents show surprising behavior in hide and seek game

September 19 2019, by Nancy Cohen

Policy Architecture



Credit: openai

Researchers have made news in letting their AI ambitions play out a formidable game of hide and seek with formidable results. The agents' environment had walls and movable boxes for a challenge where some were the hiders and others, seekers. Much happened along the way, with surprises.

Stating what was learned, the authors blogged: "We've observed [agents](#) discovering progressively more complex tool use while playing a simple game of hide-and-seek," where the agents built "a series of six distinct strategies and counterstrategies, some of which we did not know our environment supported."

In a new paper released earlier this week, the team revealed results. Their paper, "Emergent Tool Use from Multi-Agent Autocurricula," had seven authors, six of which had OpenAI representation listed, and one, Google Brain.

The authors commented on what kind of challenge they were taking on. "Creating intelligent artificial agents that can solve a wide variety of complex human-relevant tasks has been a long-standing challenge in the artificial intelligence community."

The team said that "we find that agents create a selfsupervised autocurriculum inducing multiple distinct rounds of emergent [strategy](#), many of which require sophisticated tool use and coordination."

Through hide-and-seek, (1) Seekers learned to chase hiders and hiders learned to run away (2) Hiders learned basic tool use—boxes and walls to build forts. (3) Seekers learned to use ramps to jump into hiders' shelter (4) Hiders learned to move ramps to far from where they will build their fort, and lock them in place (5) Seekers learned they can jump from locked ramps to boxes and surf the box to the hiders' shelter and (6) Hiders learned to lock the unused boxes before building their

fort.

These six strategies emerged as agents trained against each other in hide-and-seek—each new strategy created a previously nonexistent pressure for agents to progress to the next stage, without any direct incentives for agents to interact with objects or to explore. The strategies were a result of the "autocurriculum" induced by multi-agent competition and dynamics of hide-and-seek.

The authors in the [blog](#) said that they learned "it is quite often the case that agents find a way to exploit the environment you build or the physics engine in an unintended way."

What was happening was a "self-supervised emergent complexity." And this "further suggests that multi-agent co-adaptation may one day produce extremely complex and intelligent behavior." The authors similarly stated in their paper that "inducing autocurricula in physically grounded and open-ended environments could eventually enable agents to acquire an unbounded number of human-relevant skills."

Douglas Heaven, *New Scientist*, really sparked readers' interest in the way he described what happened:

"At first, the hiders simply ran away. But, they soon worked out that the quickest way to stump the seekers was to find objects in the environment to hide themselves from view, using them like a sort of tool. For example, they learned that boxes could be used to block doorways and build simple hideouts. The seekers learned that they could move a ramp around and use it to climb over walls. The bots then discovered that being a team-player—passing objects to each other or collaborating on a hideout—was the [quickest](#) way to win."

This was an ambitious project. Examining their work, *MIT Technology*

Review noted that the AI learned to use tools after nearly 500 million games of hide and seek. Through playing hide and seek hundreds of millions of rounds, two opposing teams of AI agents developed complex hiding and seeking strategies.

Karen Hao presented an interesting marker of what the agents learned after how many rounds: "...around the 25-million-game mark, play became more sophisticated. The hiders learned to move and lock the boxes and barricades in the environment to build forts around themselves so the seekers would never see them."

More millions of rounds: seekers discovered a counter-strategy, as they learned to move a ramp next to the hiders' fort and use it to climb over the walls. More rounds later, the hiders learned to lock the ramps in place before building their fort.

Yet more strategies popped up at the 380-million-game mark. two more strategies emerged. The seekers developed a strategy to break into the hiders' fort by using a locked ramp to climb onto an unlocked box, then "surf" their way on top of the box to the fort and over its walls. In the final phase, the hiders once again learned to lock all the ramps and boxes in place before building their fort.

Hao quoted Bowen Baker, one of the authors of the paper. "We didn't tell the hiders or the seekers to run near a box or interact with it...But through multiagent competition, they created new tasks for each other such that the other team had to [adapt](#)."

Think about that. Baker said they did not tell the hiders, and they did not tell the seekers, to run near boxes nor to interact with them.

Devin Coldewey in *TechCrunch* thought about it. "The study intended to, and successfully did look into the possibility of machine learning agents

learning [sophisticated](#), real-world-relevant techniques without any interference of suggestions from the researchers."

Coldewey nailed the take-home for all this work. "As the authors of the paper explain, this is kind of the way we came about."

We, as in human beings. Coldewey quoted a passage from their paper.

"The vast amount of complexity and diversity on Earth evolved due to co-evolution and competition between organisms, directed by natural selection. When a new successful strategy or mutation emerges, it changes the implicit task distribution neighboring agents need to solve and creates a new pressure for adaptation. These evolutionary arms races create implicit autocurricula whereby competing agents continually create new tasks for each other."

More information: Emergent Tool Use from Multi-Agent Autocurricula, [d4mucfpksywv.cloudfront.net/em ...
t_Emergence_2019.pdf](https://d4mucfpksywv.cloudfront.net/emergent_tool_emergence_2019.pdf)

openai.com/blog/emergent-tool-use/

© 2019 Science X Network

Citation: AI: Agents show surprising behavior in hide and seek game (2019, September 19) retrieved 19 April 2024 from <https://techxplore.com/news/2019-09-ai-agents-behavior-game.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.
